

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Methods and Systems for Preparing Extensible Markup
Language (XML) Documents and for Responding to
XML Requests**

Inventor(s):

Brian Deen
Joel Soderberg

ATTORNEY'S DOCKET NO. MS1-390US

1 **TECHNICAL FIELD**

2 This invention relates to methods and systems for preparing Extensible
3 Markup Language (XML) documents and, in particular, to methods and systems
4 for responding to XML requests.
5

6 **BACKGROUND**

7 Extensible Markup Language (XML) is a meta-markup language that
8 provides a format for describing structured data. XML is a subset of Standard
9 Generalized Markup Language (SGML) that provides a uniform method for
10 describing and exchanging structured data in an open, text-based format, and
11 delivers this data by use of standard HTTP protocol. XML utilizes the concepts of
12 elements and namespaces. XML is similar to HTML, but facilitates more precise
13 declarations of content and more meaningful search results across multiple
14 platforms.

15 XML "elements" are structural constructs that consist of a start tag, an end
16 or close tag, and the information or content that is contained between the tags. A
17 start tag is formatted as "<tag name>" and an end tag is formatted as "</tag
18 name>". In an XML document, start and end tags can be nested within other start
19 and end tags. All elements that occur within a particular element must have their
20 start and end tags occur before the end tag of that particular element. This defines
21 a strict tree-like structure that can be used to generate an XML document, or by an
22 XML parser to organize and parse the XML document. Each element forms a
23 node in this tree, and potentially has "child" or "branch" nodes. The child nodes
24 represent any XML elements that occur within the start and end tags of the
25 "parent" node.

XML accommodates an infinite number of database schemas. Within each schema, data is represented by element names. Each schema is able to define its own "dictionary" of element names, referred to as a "namespace." Namespace identifiers are used within an XML document to qualify element names, thereby allowing the same names to be used within different schemas.

Namespace inheritance within an XML document allows non-qualified names to use "default" namespaces. The default namespace for any particular XML element is whatever default namespace is applicable to the parent of the particular element. A namespace specification within an XML document is said to have a "scope" which includes all child nodes beneath the namespace specification.

Typically, XML documents are exchanged between different entities, such as client and server computers, in the form of requests and responses. A client might generate a request for information or a request for a certain server action, and a server might generate a response to the client that contains the information or confirms whether the certain action has been performed. In many cases, it is convenient to represent these XML documents in memory as a hierarchical tree structure. Once the hierarchical tree structure is built, the actual XML document in proper syntactic form can then be assembled. Consider the following exemplary XML code:

```
<orders xmlns:person="http://www.schemas.org/people"
  xmlns:dsig="http://dsig.org">
  <order>
    <sold-to>
      <person:name>
        <person:last-name>Layman</person:last-name>
        <person:first-name>Andrew</person:first-name>
      </person:name>
```

1 </sold-to>
2 <sold-on>1997-03-17</sold-on>
3 <dsig:digital-signature>1234567890</dsig:digital-
4 signature>
5 </order>
6 </orders>

7 This code includes two XML namespace declarations that are each
8 designated with "xmlns". The declarations include a prefix, e.g. "person" and
9 "dsig" respectively, and the expanded namespace to which each prefix refers, e.g.
10 "http://www.schemas.org/people", and "http://dsig.org" respectively. This code
11 tells any reader that if an element name begins with "dsig:" its meaning is defined
12 by whoever owns the "http://www.dsig.org" namespace. Similarly, elements
13 beginning with the "person:" prefix have meanings defined by the
14 "http://www.schemas.org/people" namespace.

15 Namespaces ensure that element names do not conflict, and clarify who
16 defined which term. They do not give instructions on how to process the
17 elements. Readers still need to know what the elements mean and decide how to
18 process them. Namespaces simply keep the names straight.

19 Fig. 1 shows how the structure of the above code can be represented in a
20 hierarchical tree structure. In Fig. 1, all of the elements or nodes are set out in an
21 exemplary tree that represents the XML document. Such a structure is typically
22 constructed in memory, with each node containing all data necessary for the start
23 and end tags of that node.

24 It has been typical in the past to build the entire tree structure, such as the
25 one shown in Fig. 1, before building an XML document itself. For large XML
26 documents, this can consume a great deal of memory and processor time. Thus, it
27 would be desirable to avoid this process if at all possible.

1 Accordingly, this invention arose out of concerns associated with providing
2 improved methods and systems for generating XML documents that do not require
3 or need a hierarchical tree structure to be built and stored in memory in order for
4 the XML document to be generated.

6 SUMMARY

7 Methods and systems for generating and sending XML documents and, in
8 particular, generating and sending an XML response to an XML client request are
9 described. In the described embodiment, an XML document is prepared and sent
10 to a client only a portion at a time. XML document portions are generated and
11 sent until an entire XML document is sent to the client. In a specific
12 implementation an XML response generator is provided and responds to a client
13 request without having to first build and save a hierarchical tree structure in
14 memory that represents the response. The response generator includes one or
15 more request method objects. There is one request method object for each
16 particular type of client request that might be received. Each request method
17 object knows and gathers the data that is needed to respond to its particular
18 associated client request. In addition, the request method object knows the order
19 in which the information must be provided.

20 The request method object calls an emitter object with the data that is
21 gathered by the request method object. The calls are made in a particular order
22 and ensure that the hierarchical nature of the response that is being built is
23 preserved. The emitter object translates the data that it receives into response
24 portions that are in proper XML syntactic form.

1 A body object is provided to manage a buffer. The emitter object calls the
2 body object with the properly-formatted XML response portions. The response
3 portions are placed in the buffer. When a defined buffer threshold is reached, the
4 buffered response portions are sent to the client.

6 **BRIEF DESCRIPTION OF THE DRAWINGS**

7 Fig. 1 is a diagram of a hierarchical tree structure that represent an XML
8 document.

9 Fig. 2 is a block diagram that shows an exemplary client/server architecture
10 in accordance with an embodiment of the invention.

11 Fig. 3 is a flow diagram that describes steps in a method for generating an
12 XML response.

13 Fig. 4 is a diagram that shows exemplary inventive objects and their
14 interaction with one another.

15 Fig. 5 is a computer system that is suitable for use in implementing
16 embodiments of the invention.

17 Fig. 6 is a diagram that shows a portion of a hierarchical tree structure that
18 represents a “multistatus” response that is discussed in the “Detailed Description”
19 section.

20 Fig. 7 is a block diagram of an XML response generator in accordance with
21 one embodiment of the invention.

22 Fig. 8 is a flow diagram that describes steps in a method for generating an
23 XML response in connection with the Fig. 7 XML response generator.

24

25

DETAILED DESCRIPTION

Exemplary Architecture

Before describing the various inventive methods and structures that are used in implementing the various functions described below, reference is made to Fig. 2 which shows but one example of an architecture that is suitable for use in connection with the invention. It is to be understood that the architecture that is described constitutes but one example and is not intended to limit the invention in any way.

A client is shown generally at 10 and includes a software application 12, a transport object 14, a TCP/IP module 16 and a parser 18. Although these components are shown as being grouped within the client, they can be separately organized outside of the client application 12. An exemplary client or application is one that generates requests for XML data and receives responses to its requests in the form of XML data streams that must be parsed. One specific example of an application is Microsoft's Outlook Express.

Transport object 14 can be any transport object that is used in connection with sending and receiving requests. In one specific example that transport object can be a Distributed Authoring and Versioning (WebDAV) transport object that is designed to work in connection with WebDAV requests and responses. Specific examples of these are given later in the document.

In operation, an application typically generates a request that is sent through the transport object 14 and the TCP/IP module 16 to a server 20. The server receives the request, processes it, and sends an XML response to the client.

An exemplary server 20 is Microsoft's Internet Information Service (IIS). The XML data is received into the TCP/IP module 16 and the transport object 14.

1 The transport object then begins pushing the data into the parser 18. The parser 18
2 then begins to operate on the XML data stream by parsing it and providing it to the
3 application 12. In this example, parser 18 is a so-called “push-model” parser
4 because XML data is pushed into it by the transport object 14. Aspects that are
5 described below are associated with the processing that takes place to build the
6 XML response that is sent by the server 20 to the client or application 12 in
7 response to a client request.

8 9 Overview

10 Methods and systems are provided in which an XML document is prepared
11 and sent to a client only a portion at a time. XML document portions are
12 generated and sent until an entire XML document is sent to the client. In a
13 specific implementation, various methods and systems respond to an XML client
14 request without the need to build and save in memory a hierarchical tree structure
15 that represents the entire XML response. In the described embodiment, an XML
16 response generator functions to enable a response to be prepared and sent to a
17 client in a piecewise fashion.

18 Fig. 3 shows a flow diagram that illustrates an exemplary method for
19 generating and sending an XML response that is explained in more detail below.
20 Aspects of the invention that are described below are described in the context of a
21 protocol that exists between a client and a server. In this specific protocol, a client
22 sends a request to the server and the server responds with a response. It is to be
23 understood, however, that principles of the invention have applicability outside of
24 the client/server protocol example that is given. Specifically, aspects of the
25

1 invention can be employed in any scenario where one desires to generate an XML
2 document, whether it be on the client or server side.

3 At step 100 a server, such as server 20 (Fig. 2) receives an XML client
4 request. The server 20 prepares only a portion of a response (Step 102) to the
5 client's request. The server 20 then sends the response portion (step 104) to the
6 client. The server 20 then determines whether the response preparation is
7 complete (step 106). If it is not, then step 106 loops back to step 102 and the
8 server 20 prepares another portion of the response. Accordingly, the server 20
9 repeats steps 102 and 104 until a complete response is sent to the client, at which
10 time the processing for that response ends (step 108).

11 This constitutes a highly desirable and timely improvement over past
12 methods which required that the entire XML response be built and saved in
13 memory before it was sent to the client. One of the advantages of the present
14 response processing becomes apparent in the context of very large responses that
15 must be prepared for certain client requests. An example of such a response is
16 called a "multistatus" response which is discussed in more detail below. By
17 sending a response to a client in a piecewise manner, the client can begin
18 processing the response portions (e.g. parsing the response portions and providing
19 the data to the client application 12) sooner that it could if the server had to build
20 the entire response, save it in memory, and send it out. This, in turn, translates to
21 improved processing speeds and reductions in the overhead processing that is
22 necessary to prepare and send the responses.

23 Fig. 4 shows some exemplary inventive programming structures that can be
24 utilized to generate and send the individual client response portions. Collectively,
25 these programming structures provide an XML response generator. In this

example, the response generator includes a request method object 110, an emitter object 112 and a body object 114. These objects work together to generate and send response portions to a client in a piecewise manner. In the described embodiment, there is a request method object 110 for each type of request that can be received from a client. Client request types are defined in terms of the HTTP verbs that are used in the request. When a request is received, the type of request is determined and then an appropriate request method object, e.g. object 110, is created. The request method object 110 performs data gathering functions in which it gathers the appropriate data to include in the client's response. Accordingly, request method object 110 constitutes but one example of a data-gathering mechanism. When the data is gathered by the request method object 110, a series of calls are made to the emitter object 112. The calls include the data that has been gathered and tell the emitter object 112 the information it needs to do its job. The request method object 110 also knows what element tags or nodes that it needs in its response. This information is also conveyed to the emitter object 112 in the calls that are made by the request method object 110. The emitter object 112 is then responsible for formatting the gathered data into an appropriate XML syntax for the response. Accordingly, the emitter object 112 constitutes but one example of a data-formatting mechanism.

In this example, the request method object 110 does not have to know anything about the syntax of the response that is going to be built by the emitter object 112. It only needs to know the information that is necessary for the response, e.g. the XML nodes, their organization and order within the XML response, any text values that are to be included in the response, and the like. Since there is a request method object 110 for each type of client request that can

1 be received, these objects only have to know the information or data that is
2 associated with their particular type of client request. In this example, the emitter
3 object 112 is primarily a mechanism by which the information or data is placed
4 into the correct syntactic format. Thus, the emitter object 112 does not have to do
5 any data gathering because the data and all other information it needs is provided
6 to it by the request method object 110.

7 When the emitter object 112 formats the response portions, it provides the
8 response portions to the body object 114. In this example, the body object 114 is a
9 response-sending mechanism that manages the sending function in which the
10 response portions are sent to the client. The body object 114 can also perform
11 other functions such as setting up so-called boiler plate portions of the response
12 (e.g. an XML prologue) that is to be sent. The body object 114 can also
13 accumulate response portions and send them to the client at an appropriate time.

14 Thus, the XML response generator is able to generate and send response
15 portion to a client in a piecewise fashion. This avoids having to build and save an
16 entire hierarchical tree structure that represents the response document.

17 18 **Exemplary Computer System**

19 Fig. 5 shows a general example of a computer 130 that can be used in
20 accordance with the invention. Various numbers of computers such as that shown
21 can be used in the context of a distributed computing environment that includes
22 client and server machines.

23 Computer 130 includes one or more processors or processing units 132, a
24 system memory 134, and a bus 136 that couples various system components
25 including the system memory 134 to processors 132. The bus 136 represents one

1 or more of any of several types of bus structures, including a memory bus or
2 memory controller, a peripheral bus, an accelerated graphics port, and a processor
3 or local bus using any of a variety of bus architectures. The system memory 134
4 includes read only memory (ROM) 138 and random access memory (RAM) 140.
5 A basic input/output system (BIOS) 142, containing the basic routines that help to
6 transfer information between elements within computer 130, such as during start-
7 up, is stored in ROM 138.

8 Computer 130 further includes a hard disk drive 144 for reading from and
9 writing to a hard disk (not shown), a magnetic disk drive 146 for reading from and
10 writing to a removable magnetic disk 148, and an optical disk drive 150 for
11 reading from or writing to a removable optical disk 152 such as a CD ROM or
12 other optical media. The hard disk drive 144, magnetic disk drive 146, and optical
13 disk drive 150 are connected to the bus 136 by an SCSI interface 154 or some
14 other appropriate interface. The drives and their associated computer-readable
15 media provide nonvolatile storage of computer-readable instructions, data
16 structures, program modules and other data for computer 130. Although the
17 exemplary environment described herein employs a hard disk, a removable
18 magnetic disk 148 and a removable optical disk 152, it should be appreciated by
19 those skilled in the art that other types of computer-readable media which can
20 store data that is accessible by a computer, such as magnetic cassettes, flash
21 memory cards, digital video disks, random access memories (RAMs), read only
22 memories (ROMs), and the like, may also be used in the exemplary operating
23 environment.

24 A number of program modules may be stored on the hard disk 144,
25 magnetic disk 148, optical disk 152, ROM 138, or RAM 140, including an

operating system 158, one or more application programs 160, other program modules 162, and program data 164. A user may enter commands and information into computer 130 through input devices such as a keyboard 166 and a pointing device 168. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are connected to the processing unit 132 through an interface 170 that is coupled to the bus 136. A monitor 172 or other type of display device is also connected to the bus 136 via an interface, such as a video adapter 174. In addition to the monitor, personal computers typically include other peripheral output devices (not shown) such as speakers and printers.

Computer 130 commonly operates in a networked environment using logical connections to one or more remote computers, such as a remote computer 176. The remote computer 176 may be another personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to computer 130, although only a memory storage device 178 has been illustrated in Fig. 5. The logical connections depicted in Fig. 5 include a local area network (LAN) 180 and a wide area network (WAN) 182. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet.

When used in a LAN networking environment, computer 130 is connected to the local network 180 through a network interface or adapter 184. When used in a WAN networking environment, computer 130 typically includes a modem 186 or other means for establishing communications over the wide area network 182, such as the Internet. The modem 186, which may be internal or external, is

1 connected to the bus 136 via a serial port interface 156. In a networked
2 environment, program modules depicted relative to the personal computer 130, or
3 portions thereof, may be stored in the remote memory storage device. It will be
4 appreciated that the network connections shown are exemplary and other means of
5 establishing a communications link between the computers may be used.

6 Generally, the data processors of computer 130 are programmed by means
7 of instructions stored at different times in the various computer-readable storage
8 media of the computer. Programs and operating systems are typically distributed,
9 for example, on floppy disks or CD-ROMs. From there, they are installed or
10 loaded into the secondary memory of a computer. At execution, they are loaded at
11 least partially into the computer's primary electronic memory. The invention
12 described herein includes these and other various types of computer-readable
13 storage media when such media contain instructions or programs for implementing
14 the steps described below in conjunction with a microprocessor or other data
15 processor. The invention also includes the computer itself when programmed
16 according to the methods and techniques described below.

17 For purposes of illustration, programs and other executable program
18 components such as the operating system are illustrated herein as discrete blocks,
19 although it is recognized that such programs and components reside at various
20 times in different storage components of the computer, and are executed by the
21 data processor(s) of the computer.

22 WebDAV

23 One of the areas of application for the described embodiment is in the
24 context of preparing and sending responses to client Web Distributed Authoring
25

1 and Versioning (WebDAV) requests. WebDAV is an extension to the HTTP/1.1
 2 protocol that allows clients to perform remote web content authoring operations.
 3 This extension provides a coherent set of methods, headers, request entity body
 4 formats, and response entity body formats that provide operations for properties,
 5 collections, locking and namespace operations. With respect to properties,
 6 WebDAV provides the ability to create, remove, and query information about
 7 Web pages, such as their authors, creation dates, etc. With respect to collections,
 8 WebDAV provides the ability to create sets of documents and to retrieve a
 9 hierarchical membership listing (like a directory listing in a file system). With
 10 respect to locking, WebDAV provides the ability to keep more than one person
 11 from working on a document at the same time. This prevents the "lost update
 12 problem," in which modifications are lost as first one author then another writes
 13 changes without merging the other author's changes. With respect to namespace
 14 operations, WebDAV provides the ability to instruct the server to copy and move
 15 Web resources.

16 In HTTP/1.1, method parameter information is exclusively encoded in
 17 HTTP headers. Unlike HTTP/1.1, WebDAV encodes method parameter
 18 information either in an Extensible Markup Language (XML) request entity body,
 19 or in an HTTP header. The use of XML to encode method parameters is
 20 motivated by the ability to add extra XML elements to existing structures, provide
 21 extensibility; and by XML's ability to encode information in ISO 10646 character
 22 sets, providing internationalization support. In addition to encoding method
 23 parameters, XML is used in WebDAV to encode the responses from methods,
 24 providing the extensibility and internationalization advantages of XML for method
 25 output, as well as input.

The following WebDAV HTTP methods use XML as a request and response format. The reader is assumed to have some familiarity with WebDAV HTTP methods or verbs. A brief description, however, of some pertinent WebDAV HTTP methods or verbs appears in the table immediately below:

WebDAV HTTP methods

| | |
|-----------|--|
| PROPPATCH | The PROPPATCH method processes instructions specified in the request body to set and/or remove properties defined on the resource identified by the Request-URI. |
| PROPFIND | The PROPFIND method retrieves properties defined on the resource identified by the Request-URI, if the resource does not have any internal members, or on the resource identified by the Request-URI and potentially its member resources, if the resource is a collection that has internal member URIs. |
| LOCK | A LOCK method invocation creates the lock specified by the lockinfo XML element on the Request-URI. Lock method requests SHOULD have a XML request body which contains an owner XML element for this lock request, unless this is a refresh request. The LOCK request may have a Timeout header. The LOCK is used to take out a lock of any access type. |
| UNLOCK | The UNLOCK method removes the lock identified by the lock token in the Lock-Token request header from the Request-URI, and all other resources included in the lock. |
| MOVE | The MOVE operation on a non-collection resource is the logical equivalent of a copy (COPY), followed by consistency maintenance processing, followed by a delete of the source, where all three actions are performed automatically. The consistency maintenance step allows the server to perform updates caused by the move, such as updating all URIs other than the Request-URI which identify the source resource, to point to the new destination resource. Consequently, the Destination header MUST be present on all MOVE methods and MUST follow all COPY requirements for the COPY part of the MOVE method. |
| COPY | The COPY method creates a duplicate of the source resource, identified by the Request-URI, in the destination resource, identified by the URI in the Destination header. |
| SEARCH | The SEARCH method allows queries against the different properties. |
| MKCOL | The MKCOL method is used to create a new collection. |

Multistatus Response

In one example, a client request is received and a “multistatus” response is prepared and sent to the client. A “multistatus” response is a special response body that is a text/xml or application/xml HTTP entity that contains a single XML element called “multistatus”, that contains a set of XML elements called

“response” elements. The table provided below summarizes element definitions of elements that appear in an exemplary multistatus response:

| Element Name | Namespace | Purpose | Description |
|--------------|-----------|--|---|
| multistatus | DAV | Contains multiple response messages. | The response description at the top level is used to provide a general message describing the overarching nature of the response. If this value is available an application may use it instead of presenting the individual response descriptions contained within the responses. |
| response | DAV | Holds a single response describing the effect of a method on resource and/or its properties. | A particular href MUST NOT appear more than once as the child of a response XML element under a multistatus XML element. This requirement is necessary in order to keep processing costs for a response to linear time. Essentially, this prevents having to search in order to group together all the responses by href. There are, however, no requirements regarding ordering based on href values. |
| href | DAV | Identifies the content of the element as a URI. | |
| propstat | DAV | Groups together a prop and status element that is associated with a particular href element. | The propstat XML element MUST contain one prop XML element and one status XML element. The contents of the prop XML element MUST only list the names of properties to which the result in the status element applies. |
| status | DAV | Holds a single HTTP status-line. | |
| prop | DAV | Contains properties related to a resource. | The prop XML element is a generic container for properties defined on resources. All elements inside a prop XML element MUST define properties related to the resource. No other elements may be used inside of a prop element. |

Exemplary Multistatus Response

An exemplary multistatus response in XML format is shown below. A portion of a corresponding tree structure for the exemplary multistatus response is shown in Fig. 6. Because client responses can be quite large, e.g. "multistatus" responses, having to build the entire tree structure in memory for the response, prior to generating and sending a client's response, can undesirably increase client and server response times and, necessarily, increase the complexity of the processing overhead.

09351733-072699
11 15 k2 /

```
<?xml version="1.0" ?>
<a:multistatus xmlns:a="DAV:"
xmlns:b="urn:uuid:c2f41010-65b3-11d1-a29f-00aa00c14882/" >
  <a:response>
    <a:href>http://server/folder/file1.html</a:href>
    <a:propstat>
      <a:status>HTTP/1.1 200 OK</a:status>
      <a:prop>
        <a:getcontentlength b:dt="int" >694</a:contentlength>
        <a:getcontenttype>text/html</a:getcontenttype>
      </a:prop>
    </a:propstat>
    <a:propstat>
      <a:status>HTTP/1.1 404 Resource Not Found</a:status>
      <a:prop>
        <author/>
      </a:prop>
    </a:propstat>
  </a:response>
  <a:response>
    <a:href>http://server/folder/test2.html</a:href>
    <a:propstat>
      <a:status>HTTP/1.1 200 OK</a:status>
      <a:prop>
        <a:getcontentlength b:dt="int" >1064</a:contentlength>
        <a:getcontenttype>text/html</a:getcontenttype>
      </a:prop>
    </a:propstat>
    <a:propstat>
      <a:status>HTTP/1.1 404 Resource Not Found</a:status>
      <a:prop>
        <author/>
      </a:prop>
    </a:propstat>
  </a:response>
</a:multistatus>
```

1 </a:prop>
2 </a:propstat>
3 </a:response>
4 </a:multistatus>

XML Response Generator

To reduce processing overhead complexities and increase client response efficiencies, an XML response generator is provided that builds and sends portions of a client response to a client one piece at a time. This enables the client to begin processing the response so that the data contained therein can be put to use in a more timely fashion. Accordingly, the piecewise processing and sending of the client response portions renders it unnecessary for an entire hierarchical tree structure to be built and stored in memory prior to generating the XML response.

Fig. 7 shows an exemplary implementation of an XML response generator 300 that is configured for use in connection with Microsoft's Internet Information Service (IIS). The response generator 300 includes one or more request method objects 302, 304, and 306. There is one request method object for each of the client request types that might be received. For example, for a PROPFIND request, a request method object "CpropFindRequest" is created. The XML response generator also includes an emitter object 308 and a body object 310. In this example, the request method objects 302, 304, and 306 together with the emitter object 308 constitute response-preparing mechanisms that prepare only portions of a response at a time. A buffer 312 is provided and, in this example, is associated with the body object 310. Buffer 312 buffers response portions that are received from the emitter object 308. The buffer has a defined threshold which, when satisfied, enables the body object 310 to send the buffered response portions to the client. The buffered response portions that are sent at any one time

1 constitute less than an entirety of a complete client response. Accordingly, the
2 body object 310 and buffer 312 cooperate to send the response portions that are
3 received from the emitter object 308 to the client. In this particular example, an
4 Internet Service Application Programming Interface (ISAPI) extension 314 is
5 provided and is a request-receiving mechanism that receives client XML requests.
6 ISAPI extension 314 is also responsible for returning the piece-wise generated
7 XML response data into IIS responsive to calls that the body object 310 makes to
8 the ISAPI extension.

9 Fig. 8 shows a flow diagram that describes steps in a method for responding
10 to an XML client request. When a client request is received, e.g. by the ISAPI
11 extension 314 (Fig. 7), step 400 determines the type of client request. This can be
12 determined from the request method, i.e. HTTP verb, that is used in the request.
13 In the above example, a client request in the form of a WebDAV PROPFIND
14 request is received. When the request type is determined, step 402 creates or
15 instantiates a request-type or request method object (e.g. objects 302, 304, or 306
16 in Fig. 7) for that particular request type. In this example, there is a request-type
17 object that can be instantiated for each particular request that might be received
18 from a client. So, in this example, for a PROPFIND request a
19 "CPropFindRequest" object is created.

20 Each request-type or request method object is aware of the information that
21 is necessary for its particular type of response to be generated. In addition, each
22 request type object knows the specific or defined order for the calls that it must
23 make to the emitter object 308 (Fig. 7). That is, because the client response is
24 being generated in a piecewise fashion without building the entire hierarchical tree
25 structure for the response, each request type object needs to make sure that the

order of the hierarchy is preserved in the individual response portions that are generated. For example (see Fig. 6), each “response” node can include only one “href” node. The “href” node must come before any “propstat” nodes. Each “propstat” node includes one “status” node and one “prop” node. In the past where an entire hierarchical tree structure was built in memory, nodes or elements could be added at any time and at any place in the tree because the response had not yet been formulated and sent to the client. Here, however, individual portions of the response are being generated and sent to the client so that a node or element cannot later be added to a response portion that has already been sent to the client. The way that the order of the hierarchy is preserved is by generating and sending calls to the emitter object 308 in a defined order that ensures that the hierarchical nature of the overall response is preserved. So, using the tree structure of Fig. 6 as an example, the request type object would want to make calls in an order that results in generating: the “multistatus” element, then the first “response” element, then the “href” element for the first response element, then the first “propstat” element for the first “response” element, then the “status” element for the first “propstat” element, then the “prop” element for the first “propstat” element, and then the sub-elements under the “prop” element, and so on. Since the response portions are being generated and sent in a serial, piecewise fashion, it is important to preserve the hierarchical nature of the client’s complete response. So, each request type object knows the information or data that it needs and the order of the calls it needs to make to the emitter object.

Once the request-type object has been created and gathers the information that is necessary for building a portion of the response, the request-type object calls the emitter object (step 404) and provides the information to the emitter

object. The emitter object 308 then takes the information provided to it and generates syntactically appropriate XML response portions. In doing so, the emitter object builds the XML response in a piecewise, node-by-node fashion. When the emitter object has built an XML response portion, it emits that response portion to the body object 310 (step 406). In this example, the body object 310 manages a buffer 312. The buffer 312 has a set threshold that defines how much XML data it can hold. This enables the body object 310 to accumulate or collect XML data (step 408). Step 410 determines whether the buffer threshold has been reached. If the threshold has not been reached or satisfied, step 410 loops back to step 406 which emits additional response portions to the body object. If, on the other hand, the buffer threshold has been reached, then step 412 sends the buffered XML response portions to the client. In this example, the body object 310 calls the ISAPI extension 314 which then returns the XML data to IIS. Step 414 checks to see whether the client response is complete. If it is not, then step 414 loops back to step 406 which emits XML data from the emitter object 308 to the body object 310. If the response is complete, however, then the processing for that response is over (step 416).

An advantage of the described embodiment is that processing of a client's XML request does not require building an entire hierarchical tree structure in memory prior to preparing the XML response and sending it to the client. Rather, client responses are generated in a piecewise, serial fashion. Individual response portions are prepared and sent to the client as the portions are generated. This can assist the client in beginning its processing of a response that might in some instances be quite lengthy. In addition, response processing advantages are achieved by separating functionalities into data-gathering functions that are

1 directed to gathering data that is specific to a particular client request that is
2 received, and data-formatting functions that format the data into syntactically
3 correct XML response portions.

4 Although the invention has been described in language specific to structural
5 features and/or methodological steps, it is to be understood that the invention
6 defined in the appended claims is not necessarily limited to the specific features or
7 steps described. Rather, the specific features and steps are disclosed as preferred
8 forms of implementing the claimed invention.

9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

09361782-072699